DUMPSARENA

Google Cloud Digital Leader exam

Google Cloud-Digital-Leader

Version Demo

Total Demo Questions: 44

Total Premium Questions: 441

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Digital transformation with Google Cloud	13
Topic 2, Exploring data transformation with Google Cloud	55
Topic 3, Innovating with Google Cloud AI	42
Topic 4, Modernizing infrastructure and applications with Google Cloud	173
Topic 5, Trust and security with Google Cloud	86
Topic 6, Scaling with Google Cloud operations	39
Topic 7, Driving value with Google Cloud	33
Total	441

QUESTION NO: 1

You are working with a user to set up an application in a new VPC behind a firewall and it is no-ticed that the user is concerned about data egress. Therefore, to provide assistance you want to con-figure the fewest open egress ports. Which of the following statement is correct?

A. Set up a high-priority (1000) rule that blocks all egress and a low-priority (65534) rule that allows only the appropriate ports.

B. Set up a low-priority (65534) rule that blocks all egress and a high-priority rule (1000) that allows only the appropriate ports.

C. Set up a high-priority (1000) rule to allow the appropriate ports.

D. Set up a high-priority (1000) rule that pairs both ingress and egress ports.

ANSWER: B

Explanation:

Set up a low-priority (65534) rule that blocks all egress and a high-priority rule (1000) that allows only the appropriate ports. This implements a least-privilege outbound-access policy while accounting for the implied egress rule in every Google Cloud VPC network. By default, VPC networks include an implied allow egress rule at priority 65535, allowing instances to send traffic to any destination unless a higher-priority rule applies.

Google Cloud firewall rules are evaluated by priority, where a lower numeric value has higher precedence. The specific allow rule at priority 1000 is evaluated before the broad deny rule at priority 65534, so the application can establish outbound connections only on the explicitly required protocols and ports. All other egress traffic matches the deny-all rule, which is evaluated before the implied allow egress rule at priority 65535. This approach avoids relying on the default permissive egress behavior and permits only the minimum required outbound communication.

For the rule-evaluation order and implied firewall rules, see VPC firewall rules documentation and Firewall rule evaluation order.

QUESTION NO: 2

Your organization is planning a migration of several applications to Google Cloud. Before selecting migration approaches, the organization needs to understand its existing environment and prioritize workloads according to business value, technical dependencies, and migration complexity.

Which two activities should your organization perform first? (Choose two.)

A. Inventory the existing workloads

B. Assess application dependencies

C. Decommission all on-premises systems immediately

D. Purchase commitments before estimating resource usage

E. Move every workload in a single migration event

ANSWER: A B

Explanation:

Inventory the existing workloads and assess application dependencies are appropriate initial migration activities. An inventory establishes what applications, servers, data stores, and supporting infrastructure exist. It provides the factual foundation needed to estimate scope, cost, risk, and migration effort.

Dependency assessment identifies how applications communicate with databases, services, networks, identity systems, and external components. Understanding these dependencies helps the organization group workloads into sensible migration waves and avoid moving an application before required services or data are available in Google Cloud.

Google Cloud recommends assessing workloads and dependencies as part of migration discovery and planning. See the migration to Google Cloud getting started guide and the workload assessment and discovery guidance.

QUESTION NO: 3

Your organization is assigning access to Google Cloud resources for a new application team. The organization wants to follow the principle of least privilege.

Which two actions should your organization take? (Choose two.)

A. Grant predefined roles containing only the required permissions

B. Assign access through Google Groups

C. Grant the Owner role to all application developers

D. Share one user account among all team members

E. Grant permissions directly to all users at the organization level

ANSWER: A B

Explanation:

Granting predefined roles that contain only the permissions required for a job function and assigning roles to Google Groups are both recommended ways to apply least privilege. Predefined roles group permissions around common job functions, helping administrators avoid granting broad permissions such as Owner when narrower permissions are sufficient.

Using Google Groups makes access management more efficient and auditable because access can be granted to a team or job-function group instead of individually to each employee. Administrators can then manage access by changing group membership. See the IAM security best practices and the IAM roles documentation.

QUESTION NO: 4

Your organization wants to improve the sustainability of its Google Cloud operations. The organization needs visibility into the carbon emissions associated with its cloud usage and wants to reduce waste from resources that are no longer needed.

Which two actions should your organization take? (Choose two.)

A. Use the Carbon Footprint dashboard to review emissions data

B. Remove idle resources that are no longer needed

C. Maintain unused resources to preserve maximum capacity

D. Copy all data to every available region

E. Disable cost reporting for cloud projects

ANSWER: A B

Explanation:

Use the Carbon Footprint dashboard and remove idle resources. The Carbon Footprint dashboard provides reports that estimate gross carbon emissions associated with Google Cloud usage. This enables the organization to understand emissions trends and identify projects or services that may require further review.

Removing idle resources reduces unnecessary consumption and cost. Teams should identify and delete or stop unused virtual machines, disks, IP addresses, and other cloud resources when they are no longer required. This supports both financial efficiency and the organization’s sustainability objectives.

See the Carbon Footprint dashboard overview and the Google Cloud Architecture Framework cost optimization guidance.

QUESTION NO: 5

An organization wants to use an open source library with a flexible ecosystem of tools to create and train its own machine learning models Which product or solution should the organization use1?

A. Cloud Functions

B. TensorFlow

C. Apache Beam

D. Dataflow

ANSWER: B

Explanation:

TensorFlow is the appropriate choice because it is an open-source machine learning framework designed to let organizations build, train, evaluate, and deploy their own models. It provides a broad, flexible ecosystem rather than a fixed, prebuilt modeling experience. Teams can define custom neural-network architectures, select training strategies, process input data, monitor experiments, and optimize models for a range of deployment targets. Its APIs support common machine learning workflows, including model development in Python and scalable training infrastructure.

The TensorFlow ecosystem also includes tools and libraries that support the full model lifecycle. For example, TensorFlow Datasets can help with data ingestion, TensorBoard provides visualizations for metrics and training behavior, and TensorFlow Serving supports production model serving. This flexibility makes TensorFlow well suited when an organization needs control over model design and training rather than simply consuming an existing cloud service. Google Cloud identifies TensorFlow as an end-to-end open-source platform for machine learning, with tools for building and deploying ML-powered applications. See the TensorFlow documentation and Google Cloud's TensorFlow on Google Cloud overview.

QUESTION NO: 6

Your organization stores business data in BigQuery. It needs to make the data easier to discover and govern across analytics teams while controlling access to sensitive data.

Which two Google Cloud capabilities can help meet these requirements? (Choose two.)

A. Use Dataplex Universal Catalog to manage metadata and discover data assets

B. Use BigQuery IAM permissions to control access to data resources

C. Give all employees the BigQuery Admin role

D. Copy sensitive datasets to every project that needs analytics access

E. Store data definitions only in local spreadsheets

ANSWER: A B

Explanation:

Dataplex Universal Catalog helps organizations discover, organize, and govern data assets across Google Cloud. It provides centralized metadata management and search capabilities that help data users find and understand available data.

BigQuery IAM permissions can control access to datasets, tables, views, and other BigQuery resources. Applying IAM roles according to job function supports least-privilege access to sensitive business data. See the Dataplex Universal Catalog overview and the BigQuery access control documentation.

QUESTION NO: 7

An organization wants to migrate a workload to the cloud without changing the application code or architecture Which migration path describes this approach?

A. Replatformed

B. Rehosted

C. Reimagined

D. Refactored

ANSWER: B

Explanation:

Rehosted is the appropriate migration path when an organization moves an existing workload to Google Cloud with no changes to its application code or underlying architecture. Often called a “lift-and-shift” migration, rehosting focuses on relocating the workload from its current environment to cloud infrastructure while preserving how the application is built and operates. This approach can help an organization migrate quickly, reduce data-center dependency, and begin realizing cloud benefits with limited initial engineering effort. For example, an existing application running on virtual machines can be moved to Compute Engine virtual machines with its code and architecture substantially unchanged. Google Cloud describes rehosting as moving an application to the cloud without making significant modifications, commonly as an initial migration step. After the workload is operating in Google Cloud, the organization can evaluate further modernization activities when appropriate, such as adopting managed services or modifying the application architecture. The defining characteristic for this scenario is the absence of code and architectural changes, which directly matches a rehosted migration. See the Google Cloud migration overview and guidance for choosing a Google Cloud migration path.

QUESTION NO: 8

An organization is running SQL Server on-premises and is struggling with capacity and management overhead. They want to modernize this database quickly by using Google products or services What should the organization do?

A. Migrate all SQL Server data to BigQuery

B. Refactor applications to use a cloud first database like Firestore

C. Perform a managed database migration to Cloud SQL

D. Export old tables from SQL Server to Cloud Storage

ANSWER: C

Explanation:

Perform a managed database migration to Cloud SQL is the appropriate choice because Cloud SQL for SQL Server provides a fully managed relational database service that is compatible with SQL Server workloads. It directly addresses the organization’s immediate concerns: Google Cloud operates the underlying infrastructure and handles routine operational tasks such as provisioning, patching, backups, replication, and high availability configuration. This lets the organization reduce server-capacity planning and database-management overhead while retaining a familiar SQL Server engine, minimizing the application changes needed for a rapid modernization effort.

Google Cloud Database Migration Service can support a migration from a self-managed SQL Server source to Cloud SQL for SQL Server. It is designed to simplify moving database workloads with managed migration workflows, helping organizations transition from on-premises infrastructure to a managed Cloud SQL deployment. After migration, Cloud SQL can scale to meet workload needs and provides operational controls through the Google Cloud console and APIs. See the Cloud SQL for SQL Server documentation and the Database Migration Service guide for SQL Server to Cloud SQL.

QUESTION NO: 9

How does a large hotel chain benefit from storing their customer reservation data in the cloud?

A. On-premises hardware access to transaction data

B. Real-time data transformation at scale within an on-premises database

C. Real-time business transaction accuracy at scale

D. Physical hardware access during peak demand

ANSWER: C

Explanation:

Real-time business transaction accuracy at scale is the key benefit because a hotel chain must keep reservation availability, bookings, cancellations, room assignments, and customer records synchronized across many properties and booking channels. Cloud data services can provide elastic capacity and highly available, managed transactional databases, allowing the organization to handle sharp demand changes without designing its data platform around a fixed amount of local infrastructure. This helps reservation updates be processed promptly and consistently, reducing the risk that separate systems act on stale availability information during busy periods. Google Cloud describes its managed relational database services as providing scalable and reliable database capabilities for transactional workloads. Its architecture guidance also emphasizes designing systems for reliability and scaling according to demand. For a hotel business, these characteristics support accurate, current reservation operations across a geographically distributed organization, while allowing technology teams to focus more on business processes and customer experiences instead of maintaining physical database hardware. See Cloud SQL overview and Google Cloud Architecture Framework: Reliability.

QUESTION NO: 10

Considering Different Storage and database options e.g. Cloud Datastore, Cloud SQL, Cloud Stor-age, etc. Which of the following statements is/are correct? (Select two answer)

A. Cloud DataStore and Cloud SQL have Terabytes + and Terabytes Capacity respec-tively.

B. Cloud Bigtable and Cloud Storage both have Petabytes + capacity.

C. Cloud Bigtable and Cloud Storage both have not Petabytes + capacity.

D. None of the above.

ANSWER: A B

Explanation:

“Cloud DataStore and Cloud SQL have Terabytes + and Terabytes Capacity respec-tively.” is correct as a high-level service-sizing comparison. Datastore mode is a horizontally scalable NoSQL database intended for workloads that can grow beyond ordinary single-server database sizes, while Cloud SQL is a managed relational database service whose instance storage capacity is measured in terabytes. Cloud SQL documentation specifies storage limits and configuration choices by engine and edition, reinforcing that it is suited to terabyte-scale relational workloads.

“Cloud Bigtable and Cloud Storage both have Petabytes + capacity.” is also correct. Cloud Bigtable is Google Cloud’s distributed, wide-column NoSQL database and is designed to scale to petabyte-scale data volumes with very high throughput. Cloud Storage is Google Cloud’s object storage service; its distributed architecture is intended for extremely large datasets, including petabyte-scale and larger storage needs. These capacity descriptions help distinguish the services at a conceptual level: Cloud SQL fits managed relational workloads, Datastore mode supports scalable document-style application data, Bigtable serves massive low-latency structured datasets, and Cloud Storage holds highly durable objects at very large scale. See the Cloud Bigtable overview and Cloud SQL documentation.

QUESTION NO: 11

Virtual Machine vCPU and memory usage for each of these categories can receive one of the following discounts? (Select Three Answer)

A. Military Discounts

B. Spot Instances

C. Committed-Use

D. Sustained-Use

E. Preemptible VMs

ANSWER: B C D

Explanation:

Committed-Use, Sustained-Use, and Spot Instances are the applicable pricing categories for Compute Engine virtual machine vCPU and memory usage. Committed-use discounts provide discounted pricing in exchange for a commitment to use eligible Compute Engine resources for a one- or three-year term. These commitments can be purchased for resources such as vCPUs and memory, giving organizations more predictable costs for steady, long-running workloads. Sustained-use discounts are automatically applied to eligible Compute Engine resources when they are used for a significant portion of a billing month; no advance commitment or manual enrollment is required. Spot Instances provide access to unused Compute Engine capacity at substantially discounted prices and are appropriate for fault-tolerant or interruptible workloads, because Google can stop or reclaim the instances when capacity is needed. Together, these models support distinct workload and cost-management needs: predictable baseline consumption, ongoing eligible usage, and flexible interruptible processing. Google's current Compute Engine documentation identifies Spot VMs as the current discounted interruptible VM offering, following the transition from the legacy preemptible VM model. See Spot VMs documentation and Compute Engine committed-use discounts overview.

QUESTION NO: 12

What is a defining feature of a non-relational database?

A. Repotting across multiple data sources

B. Queries that join multiple tables

C. A strictly enforced schema

D. A flexible data model

ANSWER: D

Explanation:

A flexible data model is a defining feature of non-relational databases. Rather than requiring every record to conform to one fixed, predefined table structure, many non-relational systems can store data whose fields and nesting differ between records. For example, document databases commonly represent information as JSON-like documents, allowing an application to add optional attributes or store nested objects without redesigning a single rigid schema for all data. This flexibility is particularly valuable when data is varied, changes frequently, or originates from multiple sources with different structures.

Google Cloud describes Firestore, a non-relational document database, as schema-less: documents in the same collection can contain different fields and data types. Cloud Bigtable, another non-relational database service, is designed for large-scale, sparsely populated data and supports a flexible column-family model. A flexible model does not mean data has no organization; applications can still validate, standardize, and govern stored data. It means the database does not impose the strictly uniform relational-table schema as its central data-model requirement. See Firestore data model documentation and Cloud Bigtable schema design documentation.

QUESTION NO: 13

What conditions be true if a VM interface wants to send packets to the external IP addresses of Google APIs and services using Private Google Access?

A. VM interface does not have an external IP address assigned.

B. VM interface is connected to a subnet where Private Google Access is disabled

C. Both A and B

D. None of the Above.

ANSWER: A

Explanation:

“VM interface does not have an external IP address assigned.” is a required condition for a VM to use Private Google Access when connecting to Google APIs and Google services through their external IP addresses. Private Google Access is designed to let workloads that have only internal IP addresses reach supported Google APIs and services without assigning an external IP address to the VM or routing traffic through a Cloud NAT gateway merely for that access. This supports private workload architectures while retaining controlled connectivity to Google-managed services.

In practice, this condition works together with the required subnet and network configuration: the VM’s interface must be attached to a subnet with Private Google Access enabled, and the VPC must satisfy the applicable Google API access requirements. The VM must also use an appropriate primary or alias internal source address. However, among the listed single-choice responses, the absence of an external IP address is the applicable required condition. Google’s documentation explicitly distinguishes this use case from VMs with external IP addresses, which can access Google APIs without relying on Private Google Access. See Configure Private Google Access and Private Google Access overview.

QUESTION NO: 14

Your finance team needs to receive notifications before Google Cloud spending exceeds its planned monthly amount. The team also wants to investigate the source of unexpected cost increases.

Which two actions should your organization take? (Choose two.)

A. Create a Cloud Billing budget with alert thresholds

B. Use Cloud Billing reports to analyze costs

C. Assume that a budget automatically stops all resources when the amount is reached

D. Grant all developers Billing Account Administrator access

E. Review costs only after the end of the fiscal year

ANSWER: A B

Explanation:

Creating Cloud Billing budgets with alert thresholds lets an organization receive notifications when actual or forecasted costs approach or exceed a planned amount. Budgets provide a financial governance mechanism, although they do not automatically cap usage or stop services.

Using Cloud Billing reports helps the organization investigate spending by viewing cost data over time and filtering it by dimensions such as project, service, SKU, and label. This supports identification of the resources and services contributing to unexpected increases. See the budgets and budget alerts documentation and the Cloud Billing reports documentation.

QUESTION NO: 15

A large travel company has thus far invested heavily in their technology team. There is strategic pressure on the company to focus on their core business and innovate to survive in certain geogra-phies and thrive in others. They are evaluating whether a move to Google Cloud will be good for them. Which of these reasons would be relevant for them? (choose two answer)

A. Application architecture won't be too involved because of serverless options.

B. The IT team can use managed services, for which Google manages underlying infrastructure and many maintenance tasks.

C. The IT team won't have to work on procuring and provisioning new hardware and refreshes to existing hardware.

D. Budgeting won't be an issue since the cloud takes care of billing.

ANSWER: B C

Explanation:

Using managed Google Cloud services can let the travel company spend less time on routine platform operations and more time delivering travel products, customer experiences, and market-specific innovation. For managed services, Google operates the underlying cloud infrastructure and performs service maintenance activities, reducing the operational effort required from the company's technology team. The precise division of responsibilities depends on the selected service, so choosing managed offerings is an effective way to reduce maintenance work while retaining the controls appropriate to the business. Google describes this division through its shared responsibility model.

Moving to Google Cloud also removes the need for the company to purchase, install, operate, and periodically refresh its own physical server hardware. Instead, teams can provision compute and other resources when needed, scale them with demand, and use consumption-based cloud services. This avoids long procurement cycles and the capital planning burden associated with data-center hardware refreshes. Google Cloud's infrastructure services provide the capacity on which customers run workloads, while Google is responsible for operating the physical infrastructure, as described in the Google Cloud infrastructure overview. Together, managed operations and reduced hardware ownership help the company focus investment and staff attention on its core business.

QUESTION NO: 16

You are a database manager working for a new product that will need millions of reading and writ-ing from the database, with zero downtime, key-value i.e. NoSQL features, no manual steps should be required to ensure consistency, repair data, synchronize writes and deletes, Which of the follow-ing database you choose?

A. Cloud SQL

B. Cloud BigTable

C. Cloud Spanner

D. Cloud Firestore

ANSWER: B

Explanation:

Cloud BigTable is the appropriate choice because it is Google Cloud's fully managed, wide-column NoSQL database service designed for extremely high-volume key-value workloads. It supports low-latency reads and writes at massive scale, including workloads that require millions of operations per second. Its throughput scales by adding nodes to a cluster, and clusters can be resized without application downtime, supporting changing traffic demands while maintaining availability.

For resilient deployments, Bigtable replication enables data to be replicated across clusters, including clusters in separate zones or regions. Applications can use replication and multi-cluster routing to improve availability and isolate read and write workloads. Replication is managed by the service: customers do not need to build custom processes to repair replicas or manually synchronize write and delete operations. Bigtable's architecture is therefore particularly suited to large-scale operational data such as time-series, IoT, personalization, and other high-throughput key-based access patterns. Google

documents that Bigtable is built for high throughput and low latency, with automatic replication capabilities for replicated instances. See the Cloud Bigtable overview and Cloud Bigtable replication overview.

QUESTION NO: 17

An organization wants to use BigQuery data analytics to understand their website performance, but wants to move only some data into the cloud.

Which environment should the organization use?

A. Private cloud

B. On-premises

C. Multi-cloud

D. Hybrid cloud

ANSWER: D

Explanation:

Hybrid cloud is the appropriate environment because it combines on-premises infrastructure with cloud services, allowing the organization to retain selected data in its existing environment while moving only the website-performance data needed for analysis into BigQuery. This supports a phased and targeted cloud adoption approach rather than requiring all data or workloads to be migrated. For example, the organization can export or ingest relevant web logs, event data, and analytics datasets into BigQuery, while keeping systems of record, sensitive datasets, or applications on premises. BigQuery is a fully managed Google Cloud data warehouse designed to analyze large datasets using SQL, so it can provide scalable analytics without the organization having to operate the underlying analytics infrastructure. A hybrid model also provides flexibility in how data is connected, transferred, governed, and accessed across the organization's on-premises and cloud environments. Google Cloud describes hybrid cloud as an approach that uses both on-premises resources and cloud resources, and its architecture guidance addresses connecting those environments. See Google Cloud: What is hybrid cloud? and BigQuery documentation.

QUESTION NO: 18

Your organization is preparing a disaster recovery strategy for a critical application running in Google Cloud.

Which two activities should be included in the disaster recovery strategy? (Choose two.)

A. Define recovery time objectives (RTOs)

B. Test recovery procedures regularly

C. Assume that multi-zone deployment removes the need for backups

D. Use the same recovery plan for every workload without considering business impact

E. Wait for an outage before documenting recovery responsibilities

ANSWER: A B

Explanation:

Defining recovery time objectives (RTOs) and recovery point objectives (RPOs) establishes measurable business requirements for recovering an application. RTO defines the maximum acceptable time to restore service after an interruption, while RPO defines the maximum acceptable amount of data loss measured in time.

Regularly testing recovery procedures is also essential. Tests validate that backups, replication, runbooks, access controls, and operational responsibilities work as expected before an actual disruption occurs. See the Google Cloud disaster recovery guidance.

QUESTION NO: 19

Which of the following statements describe the features of a preemptible VM in-stance? (Select Three Answer)

A. Instance is alive for no more than 12 hours

B. Can be pre-empted with a 30 minute notice

C. Can be pre-empted with a 30 second notice

D. Discounted Significantly

E. Instance is alive for no more than 24 hours

F. Can use free tier credits

ANSWER: C D E

Explanation:

Preemptible VM instances were designed for fault-tolerant, interruptible workloads that can take advantage of surplus Compute Engine capacity at a substantially reduced price. Therefore, *Discounted Significantly* is a defining characteristic: the lower price makes these instances suitable for batch processing, rendering, large-scale testing, and other jobs that can be stopped and restarted. Compute Engine can reclaim the underlying capacity when needed, so *Can be pre-empted with a 30 second notice* is also correct. The instance receives a short shutdown warning, giving software an opportunity to checkpoint work, save state, or shut down cleanly. Finally, *Instance is alive for no more than 24 hours* is correct because a preemptible VM is terminated after running for up to 24 hours, even if Compute Engine has not otherwise reclaimed its capacity. Workloads using this model should be architected to tolerate interruption and resume from saved progress. Google now recommends Spot VMs for new interruptible workloads; however, these listed behaviors describe the legacy preemptible VM offering. See the Google Cloud documentation for preemptible VMs and the Spot VM documentation.

QUESTION NO: 20

Your organization wants to run a container-based application on Google Cloud. This application is expected to increase in complexity. You have a security need for fine-grained control of traffic between the containers. You also have an operational need to exercise fine-grained control over the application's scaling policies.

What Google Cloud product or feature should your organization use?

A. Google Kubernetes Engine cluster

B. App Engine

C. Cloud Run

D. Compute Engine virtual machines

ANSWER: A

Explanation:

Google Kubernetes Engine cluster is the appropriate choice because it provides the Kubernetes orchestration capabilities needed for a containerized application that is expected to grow in complexity. GKE gives teams direct control over Kubernetes workload configuration, including Deployments, Services, autoscaling behavior, resource requests and limits, and scaling policies. For example, Horizontal Pod Autoscaling can scale workloads based on resource utilization or custom metrics, while cluster autoscaling can adjust the underlying node capacity as demand changes.

GKE also supports fine-grained traffic security and management between application components. Kubernetes NetworkPolicy can restrict which Pods may communicate with other Pods, namespaces, or IP ranges. For more advanced service-to-service controls, GKE can be used with Cloud Service Mesh, which provides a service-mesh-based approach to managing, securing, and observing traffic among services, including authorization policies and mutual TLS encryption. These capabilities make GKE suitable where teams need detailed operational and network controls rather than a more fully

managed, opinionated container runtime. See the Google Kubernetes Engine overview and GKE Network Policy documentation.

QUESTION NO: 21

You decide to migrate your on-premises environment to the cloud. You need to determine which resource components still need to be assigned ownership.

Which two functions are owned by a public cloud provider? (Choose two.)

A. Hardware maintenance

B. Infrastructure architecture

C. Infrastructure deployment automation

D. Hardware capacity management

E. Fixing application security issues

ANSWER: A D

Explanation:

Hardware maintenance and Hardware capacity management are responsibilities of a public cloud provider under the shared responsibility model. The provider operates the underlying physical environment that makes cloud services available, including data-center facilities, physical servers, storage hardware, networking equipment, and the supporting systems used to keep that hardware operational. This includes repairing or replacing failed components, performing lifecycle maintenance, and maintaining the physical infrastructure on which customer workloads run.

The provider also plans, supplies, and expands the underlying hardware fleet to meet aggregate service demand. Hardware capacity management includes procuring, installing, and operating sufficient physical compute, storage, and network capacity across the provider's cloud regions. Customers consume that capacity through cloud services rather than purchasing and maintaining their own servers or data-center equipment.

Google describes its responsibility as securing and operating the infrastructure that runs Google Cloud, while customers retain responsibilities for the resources, configurations, applications, and data they place in the cloud. This division lets organizations avoid ownership of physical hardware operations while retaining control over their cloud environment and workloads. See Google Cloud's shared responsibility and shared fate model and Google Cloud Architecture Framework guidance.

QUESTION NO: 22

An organization is concerned about the risk of data loss that may occur due to hardware failures or cyber attacks. They want to restore their systems to a previous state if such an event occurs. What should the organization do?

A. Use Cloud Monitoring.

B. Back up data regularly.

C. Set service level objectives (SLOs).

D. Enable autoscaling.

ANSWER: B

Explanation:

Back up data regularly is correct because backups create recoverable copies of data and system state that an organization can use after a hardware failure, accidental deletion, corruption, ransomware incident, or other cyberattack. A sound backup strategy defines how often copies are created, how long they are retained, where they are stored, and how recovery is

tested. Keeping backup copies isolated from primary production systems is especially important for cyber resilience, since an attack affecting the primary environment should not also compromise the copies needed for restoration.

In Google Cloud, organizations can implement backup and disaster-recovery plans with services appropriate to their workloads. For example, Cloud Storage can retain backup objects, and database services provide backup and recovery capabilities. Google Cloud also recommends planning recovery around business requirements such as recovery point objective (the maximum acceptable data loss) and recovery time objective (the acceptable time to restore service). Regularly testing restores validates that backups are complete, accessible, and usable when an actual incident occurs. Google's backup and disaster recovery guidance is available in the Backup and DR documentation, and its resilience planning concepts are covered in the Disaster recovery planning guide.

QUESTION NO: 23

Your organization wants to strengthen access security for employees who administer Google Cloud resources. The organization wants users to authenticate through its existing identity provider and to require an additional verification factor when users sign in.

Which two actions should your organization take? (Choose two.)

A. Configure single sign-on with the existing identity provider

B. Enforce multi-factor authentication

C. Assign the Owner role to all administrators

D. Create shared administrator accounts

E. Allow anonymous access to administrative applications

ANSWER: A B

Explanation:

Configure single sign-on and enforce multi-factor authentication. Single sign-on lets the organization federate authentication for its managed Google accounts with an existing identity provider. This centralizes authentication and account lifecycle management in the organization's identity system.

Multi-factor authentication requires users to provide an additional verification factor beyond their password. This reduces the risk that a compromised password alone can be used to access Google Cloud resources. Used together, federation and multi-factor authentication provide stronger identity controls for administrative access.

See the Google Workspace documentation for setting up single sign-on and enforcing 2-Step Verification.

QUESTION NO: 24

What is the benefit of using a unified cloud data solution?

A. Data will always cost less to store

B. Data can be automatically secured from external threats

C. Data can enable innovation because it's no longer siloed.

D. Data will always be backed up and cannot be lost or deleted

ANSWER: C

Explanation:

Data can enable innovation because it's no longer siloed. A unified cloud data solution brings together data that might otherwise reside in separate operational systems, analytics platforms, and business domains. By making governed data

more discoverable and accessible across an organization, teams can combine datasets, develop more complete insights, and use analytics and AI/ML capabilities to make better decisions or create new products and services. This reduces the friction of locating, preparing, and integrating data for each new use case, while helping people work from a more consistent foundation of organizational information. In Google Cloud, services such as BigQuery support a unified data-to-AI approach: data can be managed and analyzed in a scalable platform and used by multiple workloads without unnecessary copies or isolated data processes. Google describes BigQuery as an AI-ready data platform designed to unify data management, analytics, and AI, enabling organizations to derive value from their data. This ability to break down data silos and make data available for responsible cross-functional use is the central innovation benefit of a unified cloud data solution. See Google Cloud BigQuery and Modern data platform architecture on Google Cloud.

QUESTION NO: 25

An organization needs to categorize a large group of photographs using pre-trained machine learning.

Which Google Cloud product or service should the organization use?

A. Vision API

B. BigQuery ML

C. AutoML Vision

D. Looker

ANSWER: A

Explanation:

Vision API is the appropriate service because it provides Google-managed, pre-trained machine learning models for analyzing image content at scale. Its label-detection capability can identify and return descriptive labels for objects, places, activities, and general concepts found in photographs. An organization can submit images individually or process many image files asynchronously from Cloud Storage, making the service suitable for categorizing a large collection without building, training, or maintaining a custom machine learning model.

The API also provides confidence scores with detected labels, allowing the organization to establish thresholds and automation rules for how photographs are categorized. This is a managed API service, so the organization can integrate image classification into an application or workflow through REST or client libraries while Google operates the underlying model infrastructure. Google documents label detection as a Cloud Vision API feature designed to detect broad sets of categories within an image. See the Cloud Vision label detection documentation and the Cloud Vision API documentation.

QUESTION NO: 26

An organization is training a machine learning model to predict extreme weather events in their country.

How should they collect data to maximize prediction accuracy?

A. Collect all weather data evenly across all cities

B. Collect all weather data primarily from at-risk cities

C. Collect extreme weather data evenly across all cities

D. Collect extreme weather data primarily from at-risk cities

ANSWER: A

Explanation:

Collect all weather data evenly across all cities is correct because a predictive model needs training data that represents the full population and operating conditions in which it will be used. For a country-wide extreme-weather prediction service, this includes observations from every city and across ordinary as well as severe weather conditions. Broad, geographically

balanced collection captures regional differences in climate, terrain, seasonality, sensor patterns, and local weather behavior. That representation helps the model learn relationships that generalize throughout the country rather than learning patterns that apply only to a limited set of locations or only to already notable events. Including routine weather observations is also essential because the model must distinguish conditions associated with extreme events from the much more frequent non-extreme baseline. Google's machine learning guidance emphasizes that training data should be representative of the real-world data the model will encounter and that data quality, coverage, and consistency materially affect model performance. A balanced collection plan should also use consistent measurements, timestamps, and labels across cities, then reserve representative data for validation and testing. See Google Cloud's guidance on preparing data for machine learning and Google's data quality guidance.

QUESTION NO: 27

An organization is concerned about their cloud costs. They want to be informed when their spending exceeds a specific threshold, rather than waiting to see their bill at the end of the month. What should the organization do?

A. Configure a budget threshold rule and alert.

B. Pause virtual machines during non-business hours.

C. Adjust project resource quota policies.

D. Use historical cost data to predict future overspend.

ANSWER: A

Explanation:

Configure a budget threshold rule and alert. Google Cloud Billing budgets let an organization define a planned spend amount for a billing account, project, folder, or other supported scope, then configure threshold rules based on a percentage of that budget. For example, the organization can be notified when actual or forecasted costs reach 50%, 90%, or 100% of its defined budget. Notifications can be sent to billing-account users and, where configured, to Pub/Sub for automated downstream actions or integrations.

This directly meets the requirement to receive timely visibility into spending during the billing period instead of discovering an overage only after the monthly invoice is available. Budget alerts support proactive financial governance: teams can investigate unexpected consumption, identify the services driving costs, and decide whether to optimize usage or obtain additional approval before costs grow further. A budget is an alerting and monitoring mechanism, so it helps stakeholders stay informed; it does not by itself cap usage or stop resources. See the official Google Cloud Billing budgets and budget alerts documentation and programmatic notifications for budget alerts for configuration details.

QUESTION NO: 28

Your organization needs to operate a production application in Google Cloud. Operations teams need visibility into application behavior and need to be notified when service performance degrades.

Which two Google Cloud services should your organization use? (Choose two.)

A. Cloud Monitoring

B. Cloud Logging

C. Cloud Billing

D. Cloud DNS

E. Artifact Registry

ANSWER: A B

Explanation:

Cloud Monitoring provides metrics, dashboards, uptime checks, and alerting policies that help operations teams observe resource and application performance. Alerting policies can notify responders when defined conditions, such as high latency or resource utilization, are met.

Cloud Logging collects and stores log data from Google Cloud resources and applications. Teams can search and analyze logs to investigate errors, diagnose incidents, and understand application behavior. Together, these services provide core observability capabilities. See the Cloud Monitoring overview and the Cloud Logging overview.

QUESTION NO: 29

An organization is concerned that one of their applications takes too long to return a result According to Google's "Four Golden Signals" which signal measures this aspect of the applications performance?

A. Traffic

B. Saturation

C. Errors

D. Latency

ANSWER: D

Explanation:

Latency measures the time required for a system to service a request—the interval from when a request is received until the application returns a response. Therefore, when users experience an application taking too long to return a result, latency is the Golden Signal that directly quantifies that performance concern. Monitoring latency as a distribution, such as median and high-percentile response times, helps teams understand both typical user experience and the slower requests that can disproportionately affect users. Google's Site Reliability Engineering guidance identifies latency as one of the four core Golden Signals for monitoring a user-facing system. It also notes that successful and failed requests may need separate latency measurements, because an error response can be returned quickly without representing healthy application performance. Tracking latency enables organizations to establish service-level objectives and detect regressions before they create broad user impact. For additional guidance, see Google's Monitoring Distributed Systems chapter in the SRE Book and the Google Cloud documentation on service monitoring and SLOs.

QUESTION NO: 30

Which of the following are the current options for paid support in GCP? (Select Three An-swer)

A. Premier

B. Standard

C. Enhanced

D. Role

E. Premium

ANSWER: B C E

Explanation:

Google Cloud's paid Customer Care offerings are **Standard**, **Enhanced**, and **Premium**. These plans provide progressively broader technical support capabilities and response commitments for organizations operating Google Cloud workloads. Standard support supplies access to technical support for customers who need assistance beyond the no-cost Basic Support service. Enhanced support adds faster response targets and additional operational-support features for more business-critical environments. Premium support is intended for organizations with the highest support needs, offering the most comprehensive service level, including the fastest response objectives for critical-impact cases and more proactive support capabilities.

The plans are designed so that a customer can select a service level aligned with the criticality, scale, and operational requirements of its cloud environment. Google's Customer Care documentation identifies Standard, Enhanced, and Premium as its paid support service levels, while Basic Support is available without a paid support plan. See the Google Cloud Customer Care overview and the Customer Care documentation for current plan descriptions, included services, and support-case response targets.

QUESTION NO: 31

There are internal compliance requirements that demand that we do not use any APIs or services that are not backed by SLAs. Which of these are acceptable for us? (Choose two answer)

A. Alpha, Beta

B. Early Access, Preview

C. General Availability

D. Deprecated, but ensure that the SLA support period is still valid.

ANSWER: C D

Explanation:

General Availability is acceptable because Google Cloud products and features at the general-availability launch stage are intended for production use and are covered by the applicable Google Cloud Service Level Agreements. This makes them suitable where a policy requires services to have an SLA. Google describes general availability as the stage at which a product is ready for broad production use, with standard support and operational commitments.

Deprecated, but ensure that the SLA support period is still valid. is also acceptable when the particular API or service remains within its documented deprecation support period and its SLA continues to apply. Deprecation communicates that customers should plan migration away from the product or feature, but it does not necessarily terminate the service immediately. During the published support window, applicable service commitments can remain in effect. Compliance teams should verify the specific product's deprecation notice, retirement date, and applicable SLA, then complete migration before the supported period ends. Google Cloud's launch-stage definitions and deprecation guidance provide the basis for confirming these requirements.

References: Google Cloud product launch stages and Google Cloud deprecations.

QUESTION NO: 32

Your customer has a reporting tool that is only occasionally used by the leadership team. Usage of it is frequent - once a week, once a month, or once the quarter. They want to run this application in a cost-effective manner. What are the compute options available on Google Cloud which would be suitable? (Choose Two answer)

A. Cloud Run

B. Cloud App Engine Standard

C. Compute Engine

D. Kubernetes Engine

ANSWER: A B

Explanation:

Cloud Run and **Cloud App Engine Standard** are suitable because both provide managed, serverless application hosting that can minimize idle compute costs for an intermittently accessed reporting tool. Cloud Run runs containerized applications and, with its default configuration, automatically scales service instances based on incoming requests, including scaling down to zero instances when no requests are being processed. This means the customer pays primarily when the reporting

application is invoked, rather than maintaining an always-running server for weekly, monthly, or quarterly use. Google documents Cloud Run's automatic scaling behavior at Cloud Run instance autoscaling.

Cloud App Engine Standard is also a fully managed serverless platform designed for web applications. When configured with automatic scaling and no minimum idle instances, it can scale according to request volume and avoid keeping application instances running during long inactive periods. This operational model fits a leadership reporting application whose use is unpredictable or infrequent while removing the need to manage underlying infrastructure. The available scaling approaches and configuration behavior are described in the App Engine standard scaling documentation.

QUESTION NO: 33

Your organization needs to protect sensitive data stored and processed in Google Cloud. The security team wants encryption protections that are applied by default and wants the option to control encryption keys for selected workloads.

Which two statements are correct? (Choose two.)

A. Google Cloud encrypts customer data at rest by default

B. Cloud Key Management Service can manage customer-controlled encryption keys

C. Customer data is stored unencrypted unless an organization encrypts it manually

D. Encryption keys can be managed only by physical hardware administrators

E. Encryption eliminates the need to manage IAM permissions

ANSWER: A B

Explanation:

Google Cloud encrypts customer data at rest by default, and Cloud Key Management Service can be used to manage customer-controlled encryption keys. Encryption at rest is applied to customer content stored in Google Cloud services without customers needing to implement their own encryption mechanism first.

For workloads with additional key-control requirements, Cloud KMS enables organizations to create and manage cryptographic keys and use customer-managed encryption keys with supported Google Cloud services. IAM permissions and audit logging can then help control and monitor key usage.

See the default encryption at rest documentation and the Cloud Key Management Service documentation.

QUESTION NO: 34

An organization wants to transform multiple types of structured and unstructured data in the cloud from various sources. The data must be readily accessible for analysis and insights.

Which cloud data storage system should the organization use?

A. Relational database

B. Private data center

C. Data field

D. Data warehouse

E. Data lake

ANSWER: E

Explanation:

A data lake is the appropriate storage system because it centralizes data from many sources and can retain structured, semi-structured, and unstructured data at scale. Unlike a system designed around a predefined relational schema, a data lake supports ingesting diverse source formats such as transaction exports, logs, documents, media, and event streams. This makes it well suited to an organization that needs a cloud-based foundation for data transformation and broad analytical access.

On Google Cloud, a data lake architecture commonly uses scalable cloud storage for the data itself, together with services for data ingestion, transformation, governance, discovery, and analytics. Data can be kept in its original form and then prepared or transformed as needed for different analytical workloads. This enables analysts, data engineers, and data scientists to access a shared source of data for reporting, exploration, machine learning, and insight generation. Google describes a data lake as a centralized repository for storing structured and unstructured data, while its modern data lake guidance explains how it supports analytics across varied data sources. See Google Cloud: What is a data lake? and Build a data lake on Google Cloud.

QUESTION NO: 35

A customer has new applications to build that has to handle both batch data and streaming data. Which product should they choose?

A. Dataprep

B. Dataflow

C. Dataproc

D. Data Fusion

ANSWER: B

Explanation:

Dataflow is the appropriate choice because it is Google Cloud's fully managed service for building and running data-processing pipelines that support both bounded batch data and unbounded streaming data. It uses the Apache Beam programming model, which lets teams define a pipeline once and run it in batch or streaming mode with the same core concepts and transformations. This makes it especially suitable for new applications that need to ingest, transform, enrich, aggregate, and deliver data from sources such as Cloud Storage, Pub/Sub, and BigQuery without operating the underlying processing infrastructure. Dataflow automatically provisions and manages worker resources, scales capacity to meet workload demand, and provides operational features such as monitoring, autoscaling, and streaming-pipeline support. For streaming workloads, it also supports event-time processing and windowing, helping applications produce accurate results when records arrive continuously or out of order. Google describes Dataflow as a unified stream and batch data processing service, directly matching the requirement to handle both types of data. See the Dataflow overview and Google Cloud's Apache Beam programming model documentation for details.

QUESTION NO: 36

Customer Managed Encryption Keys (CMEK) can be used for encrypting data inside Cloud BigTable, which of the following statements is/are correct. (Select two answer)

A. Administrators can not rotate

B. Not supported for instances that have clustered in more than one region.

C. CMEK can only be configured at the cluster level.

D. You can not use the same CMEK key in multiple projects

ANSWER: B C

Explanation:

“Not supported for instances that have clustered in more than one region.” is correct because Cloud Bigtable CMEK has a location constraint: the Cloud KMS key used to protect a cluster must be in a location compatible with that cluster. Consequently, an instance with clusters distributed across multiple regions cannot use CMEK in this configuration. This limitation matters when designing replicated Bigtable deployments, since encryption-key requirements must be considered alongside replication and regional-resilience requirements.

“CMEK can only be configured at the cluster level.” is also correct. A Bigtable instance can contain one or more clusters, and CMEK protection is associated with each cluster rather than being a single instance-wide setting. When creating a CMEK-protected cluster, the organization selects a Cloud KMS key in an appropriate location and grants the required Bigtable service account permission to use it. This provides customer control over key rotation, access, disabling, and audit logging while Bigtable continues to manage encryption of data at rest. See the Cloud Bigtable CMEK documentation and the Cloud KMS CMEK overview.

QUESTION NO: 37

Your organization runs an application on virtual machines in Google Cloud. This application processes incoming images. This activity takes hours to create a result for each image. The workload for this application normally stays at a certain baseline level, but at regular intervals it spikes to a much greater workload. Your organization needs to control the cost to run this application.

What should your organization do?

A. Purchase committed use discounts for the baseline load

B. Purchase committed use discounts for the expected spike load

C. Leverage sustained use discounts for your virtual machines

D. Run the workload on preemptible VM instances

ANSWER: A

Explanation:

Purchase committed use discounts for the baseline load is the appropriate cost-control approach because the application has a predictable, ongoing level of virtual machine consumption. Compute Engine committed use discounts provide lower prices in exchange for a one- or three-year commitment to a specified amount of eligible resource usage. This makes them well suited to the stable portion of demand that the organization expects to run continuously, allowing it to receive a discount without committing to capacity it may not consistently need.

The regular spikes should remain outside the commitment and can be supplied by flexible, pay-as-you-go capacity as needed. Sizing the commitment to the normal baseline protects the organization from paying for unused committed resources during lower-demand periods while still reducing the cost of the usage it knows it will have. Google recommends committed use discounts for predictable, steady-state workloads and emphasizes matching commitments to anticipated resource usage. See the Compute Engine committed use discounts overview and Google Cloud’s cost optimization guidance.

QUESTION NO: 38

Google Cloud Platform (GCP) provides three main compliance resource webpages. What are they? (Select Three Answer)

A. Compliance Reports Manager

B. Support Hub

C. Compliance Offerings

D. GDPR Resource Center

E. TechCentral

ANSWER: A C D

Explanation:

Google Cloud organizes key compliance information through its Compliance Reports Manager, Compliance Offerings, and GDPR Resource Center. Compliance Reports Manager is the controlled-access location where eligible customers can request and obtain Google Cloud compliance reports and related documents, supporting their own audit, risk, and regulatory-review activities. Compliance Offerings provides an overview of Google Cloud's compliance programs, certifications, attestations, and regional or industry-focused requirements, helping organizations evaluate services against applicable obligations. The GDPR Resource Center is Google Cloud's central resource for information relevant to the General Data Protection Regulation, including Google's privacy commitments, data-processing terms, and guidance intended to support customer GDPR compliance efforts. Together, these resources address assurance documentation, the breadth of compliance coverage, and GDPR-specific privacy information. Google Cloud's compliance documentation is available at Compliance Reports Manager and Compliance Offerings; GDPR materials are available through the GDPR Resource Center.

QUESTION NO: 39

How would a global organization benefit from managing their data with Cloud Spanner?

A. Cloud Spanner is optimized for cold storage

B. Cloud Spanner replicates data across regions in real time

C. Cloud Spanner is optimized to ingest unstructured data

D. Cloud Spanner visualizes and analyzes data in real time

ANSWER: B

Explanation:

Cloud Spanner replicates data across regions in real time is correct because Cloud Spanner is a fully managed, globally distributed relational database designed for applications that need both strong transactional consistency and horizontal scale across geographic locations. A global organization can configure a regional, dual-region, or multi-region instance; multi-region configurations synchronously replicate data among geographically separated replicas. This design supports high availability and resiliency while allowing users and services in different locations to work with a single logical database. Cloud Spanner also provides external consistency for transactions, meaning committed transactions appear in a consistent order across the database. These capabilities make it suitable for globally deployed, mission-critical operational workloads, such as financial, inventory, customer, or order-management systems, where organizations need reliable access to current, consistent data across regions. Google describes Cloud Spanner as combining relational semantics with non-relational horizontal scaling and global distribution. Its replication architecture is therefore a direct benefit for an organization operating internationally, rather than merely a storage or analytics capability. See the Cloud Spanner instance configurations documentation and the Cloud Spanner overview.

QUESTION NO: 40

Which of the following is/are true about Bare Metal Solutions?

A. Enterprise-grade deployment platform

B. All your existing investment in tooling and best practices will work as is

C. Continue to run any version, and feature set, any database option, and any cus-tomizations (patchsets)

D. All of the Above.

ANSWER: D

Explanation:

All of the Above. is correct because Bare Metal Solution is designed to support specialized, enterprise workloads—particularly Oracle workloads—that need dedicated physical infrastructure while being integrated with Google Cloud services. It provides an enterprise-grade deployment environment consisting of customer-dedicated bare-metal servers, storage, and networking in Google facilities, with private, low-latency connectivity to Google Cloud. This architecture lets organizations migrate eligible Oracle applications with minimal changes and continue using established operational processes, tools, and practices.

For Oracle Database workloads, Bare Metal Solution supports a broad range of Oracle Database versions and configurations, including customer-specific customizations and patching practices, subject to the customer's licensing and support arrangements. It also supports high-availability architectures such as Oracle Real Application Clusters (RAC). Consequently, the statements concerning an enterprise-grade platform, retaining existing tooling and best practices, and continuing to operate Oracle versions, features, database options, and custom patch sets collectively describe the intended Bare Metal Solution use case. Google documents the service architecture and supported enterprise workload approach in the Bare Metal Solution overview and its guidance for Oracle workloads on Bare Metal Solution.

QUESTION NO: 41

Cloud Data Loss Prevention (DLP) is a fully managed service designed to help discover, classify, and protect the most sensitive data. DLP provides three key features (Select Three Answers)

A. Classification

B. De-identification

C. De-classification

D. Inspection

E. Reinspection

ANSWER: A B D

Explanation:

Cloud Data Loss Prevention (now part of Sensitive Data Protection) supports organizations in locating and protecting sensitive information through inspection, classification, and de-identification capabilities. **Inspection** scans supported content, such as text, images, and data stores, to detect sensitive data using built-in or custom detectors. Detection results include findings and likelihood information, enabling teams to understand where sensitive data exists. **Classification** applies context to discovered data, helping organizations identify the types and sensitivity of data they hold and prioritize governance or protection actions. It can use infoTypes, inspection templates, and data profiles to support data discovery and classification across Google Cloud environments. **De-identification** protects detected sensitive values by transforming them, for example through masking, tokenization, bucketing, date shifting, or encryption-based transformations. This allows data to remain useful for analytics, testing, or sharing while reducing exposure of direct identifiers. Together, these capabilities support the lifecycle described in the question: discovering sensitive information, understanding its sensitivity, and applying protective transformations. See the Sensitive Data Protection inspection overview and de-identification documentation.

QUESTION NO: 42

You're negotiating SLAs with a customer. You have communicated that there will be a 99.99% (four 9s) availability for the service you are providing. Every aspect of the service is under your con-trol. They want to modify the reliability to 99.999% (five 9s). What do you tell them? (Choose two answer)

A. Yes, that could be possible. If yes, there will be a significantly higher charge be-cause the effort is significantly higher too.

B. Yes, that is possible, but there will be an additional charge of 9% for the service because that is the additional effort required.

C. Yes, that is possible. There is hardly any difference to provide another 0.009% availability.

D. Ask them for the reasonable downtime they are willing to absorb. If it is more than 60 minutes in an entire year, explain how the current SLA meets that requirement.

ANSWER: A D

Explanation:

“Yes, that could be possible. If yes, there will be a significantly higher charge be-cause the effort is significantly higher too.” is appropriate because a fifth nine is not a marginal operational commitment. Over a full year, 99.99% availability permits about 52.56 minutes of downtime, whereas 99.999% permits only about 5.26 minutes. Achieving that target generally demands substantially stronger architecture and operations: eliminating single points of failure, using resilient multi-zone or multi-region designs where justified, automating recovery, continuously monitoring service health, and rigorously testing failure scenarios. These capabilities add engineering, infrastructure, and operational cost, so the SLA and price should reflect the increased commitment.

“Ask them for the reasonable downtime they are willing to absorb. If it is more than 60 minutes in an entire year, explain how the current SLA meets that requirement.” is also correct because an SLA should be based on the customer’s business requirements and downtime tolerance, rather than selecting a larger number of nines without understanding its value. A four-nines annual target is already within roughly one hour of downtime, so it meets a requirement that tolerates more than 60 minutes annually. Google recommends defining reliability targets from user and business needs, then balancing reliability investments against cost and complexity. See the Google Cloud Architecture Framework: Reliability and Google SRE Workbook: Implementing SLOs.

QUESTION NO: 43

Your organization is developing a plan for migrating to Google Cloud.

What is a best practice when initially configuring your Google Cloud environment?

A. Create a project via Google Cloud Console per department in your company

B. Define your resource hierarchy with an organization node on top

C. Create projects based on team members’ requests

D. Make every member of your company the project owner

ANSWER: B

Explanation:

Define your resource hierarchy with an organization node on top is the recommended best practice because the organization resource is the root node of a Google Cloud resource hierarchy. It provides a central, company-wide boundary for governing cloud resources and enables consistent administration as the environment grows. Under the organization node, an enterprise can create folders to represent major business units, environments, or other logical groupings, then place projects within those folders. This structure supports the inheritance of organization policies and IAM access controls downward through folders and projects, helping teams apply governance, security, and compliance requirements consistently rather than managing each project in isolation.

Establishing this hierarchy early is particularly important during a migration because it creates a scalable foundation for billing, identity, access management, policy enforcement, and resource organization before workloads are deployed. Google Cloud recommends using the organization resource to centrally manage projects and resources associated with a company’s Google Workspace or Cloud Identity account. A thoughtfully designed hierarchy also lets administrators delegate appropriate responsibilities at folder or project levels while retaining centralized oversight at the organization level. See the Google Cloud resource hierarchy documentation and Google Cloud’s enterprise organization best practices.

QUESTION NO: 44

An organization decides to migrate their on-prenwses environment to the cloud They need to determine which resource components still need to be assigned ownership

Which two functions does a public cloud provider own? (Choose 2)

Choose 2 answers

A. Fixing application security issues

B. Infrastructure architecture

C. Hardware capacity management

D. Hardware maintenance

E. Infrastructure deployment automation

ANSWER: C D

Explanation:

Hardware capacity management and **Hardware maintenance** are responsibilities of a public cloud provider under the cloud shared-responsibility model. The provider operates the underlying physical infrastructure that delivers cloud services, including data centers, servers, storage devices, networking equipment, power, cooling, and the facilities that house those components. This includes monitoring and planning the physical resource capacity needed to run the provider's cloud platform, replacing failed components, and maintaining the hardware fleet.

For a customer using public cloud services, this removes the need to purchase, install, refresh, or repair physical servers and associated data-center equipment. Instead, the organization can consume infrastructure resources on demand and focus its ownership responsibilities on the workloads and configurations it deploys in the cloud. Google Cloud describes this division of responsibilities as Google securing the cloud infrastructure while customers retain responsibility for what they run and configure in the cloud. See Google Cloud's shared responsibility model and Google Cloud overview documentation.